

[0001] SOURCE SYNCHRONOUS LINK WITH CLOCK
RECOVERY AND BIT SKEW ALIGNMENT

[0002] BACKGROUND OF THE INVENTION

A. Field of the Invention

[0003] The present invention relates generally to high-speed data transmission, and more particularly, to a source synchronous data transmission link.

B. Description of Related Art

[0004] Conventional communication systems typically use either synchronous or asynchronous data transmission to transmit data from a transmitting system to a receiving system. Synchronous data transmission requires the transmitting and receiving systems to be synchronized with each other.

[0005] Conventional synchronous data transfer has the characteristic of requiring relatively constant time periods between successive events. Keeping time periods relatively constant, such as the time between successive bits, allows end systems to stay synchronized because the timing of successive events is predetermined.

[0006] To implement synchronous data transfer, the transmitting and receiving systems can be synchronized in a variety of ways. In some synchronous data transmission systems, for example, the transmitting and receiving ends share a common master clock. Because the end systems are

synchronized with each other, synchronizing information does not need to be sent with the data. Although synchronous data transfer has many advantages, the precise clocks and timing requirements are very costly.

[0007] For asynchronous data transfer, data can be transmitted at irregular intervals. There is no clocking signal in asynchronous data transfer.

Transferring data at irregular intervals, however, requires insertion of start and stop bits to mark the beginning and end, respectively, of the data stream. The receiving system uses the start and stop bits to determine when data is being received. Thus, although asynchronous data transfer is less costly in terms of clocks and maintaining timing, it lowers bandwidth because of the overhead of start and stops bits in the transmitted data. Therefore, conventional synchronous and asynchronous data transmission systems each have disadvantageous.

[0008] Accordingly, it is desirable to efficiently transmit data signals between devices without the disadvantages of conventional synchronous or asynchronous data transmission systems.

[0009] SUMMARY OF THE INVENTION

[0010] Systems and methods consistent with the invention address, among other things, the need to efficiently transmit data by improving the density of links connecting a plurality of signals.

[0011] In one aspect, the principles of the invention are directed to a communication method that includes receiving data from data lines, each data

line providing data at a predetermined rate, and serializing the received data.

Additionally, the serialized data is provided through a optical link and deserialized. The deserialized data is provided to second data lines corresponding to the first data lines.

[0012] A second aspect of the principles of the present invention is directed to a communication system that includes input data signal lines carrying data transmitted synchronously with a first clock signal running at a predetermined frequency. A plurality of multiplexers are each configured to receive a subset of the input data signal lines and to combine data received on the subset of data signal lines into combined signals. At least one of the multiplexers receiving, as the subset of the data signal lines, a first of the data signal lines and a compliment of the first of the data signal lines. Demultiplexers convert the combined signals to data signals on output data signal lines based on a second clock signal, and a clock recovery circuit provide the demultiplexers with the second clock signal based on the combined signals.

[0013] Another aspect of the principles of the present invention is directed to a network device for transmitting information over a link. The network device comprises a first multiplexer connected to data lines, the first multiplexer combining data signals from the data lines and a second multiplexer including a first input connected to the data lines, a second input connected to a compliment of the first input, and an output, the second multiplexer combining data signals from the first and second inputs. Further, transmitters are connected to the outputs of the first and second multiplexers, the transmitters receive the

combined data signals from the first and second multiplexers and transmit the combined data signals to a link.

[0014] Another network device consistent with the principles of the present invention receives information. Receivers are connected to a link, the receivers receiving the information from the link and converting the information to first data signals and transmitting the data signals on data lines. A first demultiplexer is connected to the data lines, the demultiplexer converting the first data signals into second and third data signals, and outputting the second data signal based on a clock signal and outputting the third data signal based on the clock signal. A clock recovery circuit is connected to the data lines, the clock recovery circuit generating the clock signal based on the first data signals and providing the clock signal to the demultiplexer. Finally, a receiver is connected to receive the second and third data signals, the receiver analyzing the second and third data signals for the presence of a predetermined pattern and using the predetermined pattern to align bits in the second and third data signals.

[0015] Additional aspects of the principles of the invention are evident from the following description and the claims.

[0016] BRIEF DESCRIPTION OF THE DRAWINGS

[0017] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate the principles of the invention and, together with the description, explain principles of the invention. In the drawings,

[0018] Fig. 1 is a block diagram illustrating a data transmission system consistent with the principles of the invention;

[0019] Fig. 2 is a diagram illustrating an example of input data links from Fig. 1;

[0020] Fig. 3 is diagram illustrating examples of portions of the electrical-optical data transmission system of Fig. 1;

[0021] Fig. 4 is a diagram illustrating an exemplary alignment pattern for aligning received data bits consistent with the principles of the invention; and

[0022] Fig. 5 is a flow chart illustrating methods associated with the alignment pattern shown in Fig. 4.

[0023] DETAILED DESCRIPTION

[0024] The following detailed description of the invention refers to the accompanying drawings. The same reference numbers are sometimes used in different drawings to identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and equivalents of the claim limitations.

[0025] As described herein, a transmission system converts a plurality of synchronous signals for transmission onto a physical link. In one embodiment, transmission bandwidth is increased by embedding a clock signal in the converted signals.

[0026] Fig. 1 is a block diagram illustrating a data transmission system 100 consistent with the principles of the invention. Data transmission system 100 includes N input synchronous data links 101-103, a multiplexing unit 110, a link 112, a demultiplexing unit 120, and N output synchronous data links 104-106. Each data link 101-103 may carry data signals transmitted synchronously with a clock signal. Multiplexing unit 110 combines data streams from data links 101-103 into single streams and transmits the single streams over link 112. In one embodiment, link 112 is an optical link containing a number of optical fibers. Demultiplexing unit 120 demultiplexes the single streams received from optical link 112 and outputs the demultiplexed data on output data links 104-106.

[0027] Fig. 2 is a diagram illustrating an example of data bit streams and clock signals for one of data links 101-103. In the illustrated embodiment, there are four data bit streams 201-204 and a clock signal 210. Data bit streams 201-204 are transmitted synchronously with clock signal 210. In the illustrated embodiment, for each clock cycle, bit streams 201-204 together define an A-bit wide data word, where A in the embodiment of Fig. 2 is four. In an embodiment having three data links ($N=3$), there would be twelve ($N \bullet A$) bits of information transmitted per clock cycle.

[0028] In one embodiment, clock signal 210 of each of data links 101-103 is an independent clock signal running at the same frequency as the clock signals (or a multiple of the frequency of the clock signals) for the other data links. Each clock signal 210 may, however, be an independent clock signal in which the phase of each clock relative to each other may vary.

[0029] Returning to Fig. 1, multiplexing unit 110 may perform a number of functions, including combining multiple data bit streams 201-204 and transmitting the combined streams on link 112. Clock signal 210 is not transmitted over link 112. Instead, multiplexing unit 110 embeds an indication of clock signal 210 in the transmitted data bits. Some bits may be scrambled and have a probabilistic embedding of the clock signal that is preset in data bit streams 201-204.

Demultiplexing unit 120 recovers the clock signal and demultiplexes the combined data streams into output data links 104-106. By not sending clock signal 112 over link 112, clock skew which may have otherwise been introduced by link 112 is avoided and fewer pins may be used.

[0030] In one embodiment, multiplexing unit 110 is implemented within a first physical chassis, such as a router chassis, and demultiplexer unit 120 is implemented within a second physical chassis, such as a second router chassis. In this embodiment, link 112 connects the two router chassis.

[0031] Multiplexing unit 110, demultiplexing unit 120, and link 112 will now be described with reference to Fig. 3. In Fig. 3, two synchronous data links are illustrated, each eleven bit streams wide (i.e., $A=11$), labeled as bit lines zero through ten. The description of Fig. 3 describes the upper input synchronous data link. The lower synchronous data link operates in a similar manner.

[0032] High speed transmitter (HST) 301 receives bit streams from a data link, such as data link 101, processes the bit streams, and forwards the bit streams to multiplexers 305-310. The bit streams are combined by multiplexers 305-310, converted into optical signals by optical transmitters 325, transmitted

over optical fibers 330, converted back to an electrical signal by optical receivers 335, and processed and reassembled into the original data by demultiplexers 340-345 and high speed receiver (HSR) 351.

[0033] Multiplexing unit 110 and demultiplexing unit 120 may additionally include corresponding scramblers (not shown) and descramblers (not shown), respectively, to scramble data transmitted over fibers 330.

[0034] Each of multiplexers 305-309 receives from HST 301 two bit streams of the eleven incoming bit streams and combines the two bit streams into combined bit streams 315-319. In one embodiment, multiplexers 305-309 interleave each input bit stream. For example, bit stream 315 contains alternating bits from bit stream number zero and bit stream number one, and is clocked at twice the frequency of the signals entering HST 301. For example, the bit streams received by HST 301 and multiplexers 305-310 may be clocked at 937 million bits per second (Mbps) while the data transmitted from multiplexers 305-310 may be clocked at 1874 Mbps.

[0035] Multiplexer 310 operates slightly differently than multiplexers 305-309. Multiplexer 310 receives only one bit stream directly, bit stream 10. The input to multiplexer 310 includes an inverter 311. Accordingly, bit stream 320 alternates between the bit value of bit stream 10 and its complement. Because each bit is followed by its complement, bit stream 320 is guaranteed to include an edge every other clock signal.

[0035] Although bit streams 315-320 are illustrated in Fig. 3 as a multiplexed stream made from two constituent streams, in alternate implementations three or more streams may be combined into a single stream.

[0036] The optical transceiving system that includes optical transmitters 325, optical fibers 330, and optical receivers 335 transmit the combined data streams 315-320 to demultiplexing unit 120. More particularly, optical transmitters 325 receive the combined data streams 315 through 320 and convert the electrical signals into optical signals for transmission over optical fibers 330. The optical signals, after traveling over the optical fibers 330, are converted back to electrical signals by optical receivers 335.

[0037] The clock signal is not transmitted over optical fibers 330. Instead, clock recovery circuit 360 derives a clock signal based on combined stream 320 after it is output by the optical receiver 335 carrying combined stream 320. Because each bit in signal 320 is followed by its complement, the signal includes an edge every other clock cycle. Clock recovery circuit 360 uses the presence of the regularly repeating edge to determine the clock frequency and the phase of the clock for the transmitted data. More particularly, clock recovery circuit 360 derives from combined stream 320 a clock out signal having the same frequency as the transmitting clock and in phase with combined data stream 320.

[0038] The clock out signal may be shared by both HSR 351 and the lower HSR shown in Fig. 3. The clock out signal generated by clock recovery circuit 360 is input to demultiplexers 340-345, which receive the combined data streams 315-320 from optical receivers 335. Based on the clock out signal,

demultiplexers 340-345 separate the combined data streams into the constituent bit streams zero through ten, and the complement of bit stream 10. The deserialization process may cause pairs of data bits, such as 0/1, 2/3, and 4/5, to become transposed.

[0039] Although clock recovery circuit 360 aligns the phase of the clock signal at the receiving end to the input data, the bit streams transmitted between HST 301 and HSR 351 may become misaligned with one another because, for example, the fiber may have different lengths and transmission properties. To address this potential problem, and to correct bit transposition, HST 301 and HSR 351 may periodically transmit preestablished alignment patterns that both HST 301 and HSR 351 recognize. Based on the received pattern, HSR 351 can check to ensure that all of its received bit streams are aligned and realign bit streams if necessary.

[0040] Fig. 4 is a diagram illustrating an exemplary alignment pattern 401 that may be used by HST 301 and HSR 351 consistent with the principles of the invention. Pattern 401 may have the property of being easily detected in the presence of skew and bit transposition. The pattern includes for each bit stream in the data link a number of consecutive bits 405 (e.g., 6) followed by an equal number of consecutive complementary bits 406. In one embodiment the pattern alternates between consecutive bit streams. Thus, as shown in Fig. 4, the pattern for bit streams 0, 2, 4, 6, 8, and 10 is a series of logic ones followed by a series of logic zeroes, while the pattern for bit streams 1, 3, 5, 7, and 9 is a series of logic zeroes followed by a series of logic ones. After the second series of bits,

[0041] In operation, HSR 351 monitors the incoming bit streams for any pattern which could have been a transmitted pattern 401 that was altered by bit skew or transposition of bits. Upon detecting the pattern, HSR 351 determines whether bit streams 0-10 are aligned. If any of the streams are out of alignment, such as if bit stream zero is shifted relative to bit stream one, HSR 351 adjusts bit stream state buffers within HSR 351 to align the bits to form pattern 401. To this end, HSR 351 may store recovered bits from each line in a FIFO (first-in-first-out) queue, each FIFO having a separate read pointer and having the ability to transpose adjacent bits.

[0042] Fig. 5 is a flow chart illustrating operation of HSR 351 in processing alignment pattern 401. HSR 351 may be processing when a CRC error burst occurs that indicates an out-of-alignment condition. HSR 351 begins by receiving the incoming bit streams and analyzing the received bit streams for the presence of pattern 401. (Acts 501 and 502). Different ones of the bit streams may be received as shifted versions of the pattern. Accordingly, HSR 351 takes this fact into account when detecting the pattern 401. When the pattern is located, (Act 503), HSR 351 adjusts receive buffers within HSR 351 (not illustrated in Fig. 4) to compensate for any shift or transposition in the bit streams. (Act 504). Thus, for example, if the second bit stream is a cycle behind all of the other bit streams, HSR 351 may delete a bit in the buffer or double clock the buffer corresponding to the second bit stream to bring the second bit

stream in alignment with the other bit streams. Acts 501 through 504 are repeated whenever the number of consecutive errors exceeds a threshold. (Acts 505 and 506). Errors may be detected through the CRC field.

[0043] As described above, consistent with the principles of the invention, data from a plurality of synchronous data links are serialized before being transmitted. A clock signal is not transmitted with the serialized data. Instead, a clock signal is generated at the receiving end of the optical link based on frequency and phase information embedded in the transmitted data.

[0044] The foregoing description of embodiments of the invention provides illustration and description, but is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the invention. Although described as being primarily implemented in hardware, the present invention may be embodied in hardware and/or in software (including firmware, resident software, micro-code, etc.).

[0045] The scope of the invention is defined by the claims and their equivalents.